

WHAT IS CLAIMED IS:

1. A method for allocating application processing operations among information handling system cluster resources in response to a fail-over event, comprising:

5 identifying a performance ratio between a failing-over cluster node and a fail-over cluster node; transforming a first calendar schedule associated with failing-over application processing operations into a second calendar schedule to be associated with failing-over application processing operations on the fail-over cluster node in accordance with the performance ratio; and

10 implementing the second calendar schedule on the fail-over cluster node such that the fail-over cluster node may effect failing-over application processing operations according to the second calendar schedule.

2. The method of Claim 1, further comprising determining whether resources on the fail-over cluster node are sufficient to support failing-over application processing operations in accordance with the second calendar schedule in addition to any existing fail-over cluster node application processing operations.

3. The method of Claim 2, further comprising
applying a resource negotiation algorithm to the
application processing operations of the fail-over node
in response to determining that the resources of the
5 fail-over cluster node are insufficient to support both
failing-over application processing operations in
accordance with the second calendar schedule and any
existing fail-over cluster node application processing
operations.

10

4. The method of Claim 3, further comprising:
calculating a new calendar schedule for the
fail-over node application processing operations based on
results from application of the resource negotiation
15 algorithm; and
implementing the new calendar schedule on the
fail-over node.

20

5. The method of Claim 1, further comprising:
identifying at least one characteristic of the
failing-over cluster node;
identifying at least one characteristic of the
fail-over cluster node; and
calculating the performance ratio between the
25 failing-over cluster node and the fail-over cluster node
based on the identified characteristics of each node.

6. The method of Claim 1, further comprising
collecting information handling system cluster node
resources required by at least one application to be
deployed in an information handling system cluster
5 configuration.

7. The method of Claim 1, further comprising
maintaining a knowledge-base containing information
regarding one or more operational aspects of the
10 information handling system cluster.

8. The method of Claim 7, further comprising
determining whether the first calendar schedule for a
selected cluster node is feasible using operational
15 aspects of the selected cluster node available in the
knowledge-base.

9. The method of Claim 1, further comprising
updating an application-to-cluster node map identifying
20 the cluster node associated with each application
following the allocation of application processing
operations among the information handling system
resources in response to a fail-over event.

10. A system for maintaining resource availability in response to a fail-over event, comprising:

an information handling system cluster including a plurality of nodes;

5 at least one storage device operably coupled to the cluster; and

a program of instructions storable in a memory and executable in a processor of at least one node, the program of instructions operable to identify at least one
10 characteristic of a failing node and at least one characteristic of a fail-over node, calculate a performance ratio between the failing node and the fail-over node, transform a processing schedule for at least one failing-over application to a new processing
15 schedule associated with failing-over application processing on the fail-over node in accordance with the performance ratio and implement the new processing schedule for the failing-over application on the fail-over node.

20 11. The system of Claim 10, further comprising the program of instructions operable to gather node resource requirements for at least one application to be deployed in the cluster.

25 12. The system of Claim 11, further comprising the program of instructions operable to gather resources available on at least one node of the cluster.

13. The system of Claim 12, further comprising the program of instructions operable to verify that the resources of a selected node are sufficient to perform processing operations in accordance with the resource
5 requirements of at least one application to be deployed on the selected node.

14. The system of Claim 10, further comprising the program of instructions operable to:
10 evaluate application processing resources available on the fail-over node; and
determine whether the application resources available on the fail-over node are sufficient to perform processing operations for the failing-over application in
15 accordance with the new processing schedule and any existing fail-over application processing operations.

15. The system of Claim 14, further comprising the program of instructions operable to:
20 apply a resource negotiation algorithm to at least the new processing schedule in response to a determination that the application processing resources of the fail-over node are insufficient to support both the processing schedule of the failing-over application
25 and any existing fail-over applications;
calculate at least one modified processing schedule in accordance with results of the resource negotiation algorithm; and
implement the modified processing schedule on the
30 fail-over node.

16. The system of Claim 15, further comprising the
program of instructions operable to apply the resource
negotiation algorithm to the new processing schedule for
the failing-over application and at least one existing
5 fail-over node processing schedule.

17. Software for allocating information handling
system resources in a cluster in response to a fail-over
event, the software embodied in computer readable media
5 and when executed operable to:

access a knowledge-base containing application
resource requirements and available cluster node
resources;

calculate a performance ratio between a failing node
10 and a fail-over node;

develop a new processing schedule for a failing-over
application on the fail-over node in accordance with the
performance ratio; and

queue the failing-over application for processing on
15 the fail-over node in accordance with the new processing
schedule.

18. The software of Claim 17, further operable to:
gather resource requirements for each application in
20 the cluster selected for fail-over protection; and
store the application resource requirements in a
static data portion of the knowledge-base.

19. The software of Claim 18, further operable to:
25 gather available resource information for each
cluster node selected for operation as a fail-over node;
and
store the available node resource information in the
static data portion of the knowledge-base.

30

20. The software of Claim 19 further operable to
determine whether a selected node includes resources
available to support a processing schedule for a selected
application based on the resource requirements of the
5 application and the available resources on the node from
information maintained in the knowledge-base.

21. The software of Claim 17, further operable to
determine whether the new processing schedule may be
10 supported by the fail-over node.

22. The software of Claim 21, further operable to:
apply a resource negotiation algorithm to each
processing schedule associated with the fail-over node;
15 generate new processing schedules for applications
to be executed by the fail-over node; and
queue the applications to be executed by the
fail-over node in accordance with resource negotiation
algorithm generated processing schedules.

20 23. The software of Claim 17, further operable to
update an application-to-node map contained in the
knowledge-base.